SRI International

AD-A257 409

Annual Progress Report - 12 October 1992

Distributed Reasoning and Planning

Contract No. N00014-89-C-0095
SRI Project ECU 7363
Covering the period 1 October 1991 through 30 September1992

Prepared by:
Dr. Kurt G. Konolige, Sr. Computer Scientist
Artificial Intelligence Center

Prepared for:
Chief of Naval Research
Code 1133/Annual Report
Ballston Tower One
800 North Quincy Street
Arlington, Virginia 22217-5660

OCT 23 1992

Approved by:
Dr. C. Raymond Perrault, Director
Artificial Intelligence Center

DEFENSE TECHNICAL INFORMATION CENTER

9227623

ONR CONTRACT INFORMATION

Contract Title:	DISTRIBUTED REASONING AND PLANNING
Contract Number:	N00014–89–C–0095
Scientific Officer:	LtCdr Robert Powell
Principal Investigator:	Kurt Konolige
Institution:	SRI International
Telephone:	(415) 859-2788
Electronic mail:	konolige@ai.sri.com
Period of Performance:	1 Oct 91 – 30 Sep 92

Accesion For

NTIS CRA&I	X
DTIC TAB	☐
Unannounced	☐

per LtCdr R. Powell

By ....

Distribution /

Availability Codes

| Dist | Avail and / or Special |
| A-1 | |

Statement A per telecom
Lt Cdr Robert Powell
ONR/Code 113D, Arlington, VA
22217-5000
NWW 10/23/92

DTIC QUALITY INSPECTED 1

Kurt Konolige
SRI International
(415) 859-2768
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 – 30 Sep 92

List of numerical productivity measures

Refereed papers submitted but not yet published:	3
Refereed papers published:	6
Unrefereed reports and articles:	5
Books or parts thereof submitted but not yet published:	1
Books or parts thereof published:	0
Patents filed but not yet granted:	0
Patents granted:	0
Invited presentations:	5
Contributed presentations:	8
Honors received:	17
Prizes or awards received:	1
Promotions obtained:	0
Graduate students supported:	0
Post-docs supported:	1
Minorities supported:	0

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 - 30 Sep 92

3

Summary of technical progress

For complex tasks, it is useful to view the components of distributed systems as specialized intelligent agents, each planning in cooperation with the others to achieve certain goals. This project involves research in the reasoning abilities of the individual agents and in interagent coordination and communication strategies. In particular, the following research areas have been or are currently being addressed.

- The architectural design of an autonomous agent acting in a dynamic environment, especially the integration of reactive and strategic planning.

- Reasoning with analogical representations, particularly with respect to using and learning maps.

- Models of intention and belief, and computations methods for reasoning about them.

Architectural design of agents. The individual agents that constitute a distributed system must be designed with computational resource limitations in mind, if they are to successfully synchronize and coordinate their activities. In previous years we developed and implemented a theory of planning and deliberation for resource-bounded agents, which we call IRMA: the Intelligent, Resource-Bounded Machine Architecture. We also built a testbed, the Tileworld, for evaluating this and similar architectures.

The IRMA architecture and the Tileworld testbed are useful for exploring trade-offs in commitment to current plans vs. deliberation about new opportunities. This is one problem real planning agents face; another is the uncertainty associated with dynamic worlds, which prevents agents from forming detailed plans *ab initio*. Tileworld assumes agents have complete knowledge of their surroundings, but this is an idealization. Real agents must react to their surroundings as the perceive them, with virtually no time to deliberate and form plans. Instead, they employ *reactive strategies*: procedures that keep them in balance with their changing world. For example, mobile agents employ such strategies to avoid obstacles as they appear in their path.

The major research issue in reactive strategies is how to integrate multiple competing goals: the immediate local goal of avoiding trouble, and the more global goal of achieving a desired position or state. For example, mobile agents navigating in a dynamic environment must avoid obstacles, but at the same time must try to do maintain their pursuit of a desired position.

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 - 30 Sep 92

Our approach to this problem is a paradigm we call *monitored behaviors*. The basic idea is that reactive strategies are structured as behaviors that incorporate a goal, and a means of deciding whether that goal is being achieved or if the behavior is *frustrated*. For example, a mobile robot may have a behavior that follows along a wall; if the robot loses track of the wall, or has to deviate too far from it, the behavior is frustrated and a monitor is notified. The monitor is responsible for either correcting the behavior or notifying the planning and deliberation processes that the desired behavior is not achievable, and replanning is necessary.

To test our theory, we have implemented a preliminary version on our mobile robot testbed, Flakey. Flakey participated in the Robotics Competition at the AAAI Conference in San Jose, CA (July 1992). We have described the theory, and Flakey's performance in the competition, in an article to appear in AAAI magazine ("CARMEL vs. Flakey: A comparison of two winners"). Flakey was the only robot to finish the competition using the "natural" environment, that is, not using any modifications to the surroundings to improve its performance. Even though this made the competition much harder, Flakey still came in second, thus validating the theory experimentally.

Analogical representations

We have initiated a project that addresses the problem of using analogical representations effectively in automated reasoning systems. Analogical representations have the property that their structure embeds properties of the domain being modeled. Maps provide a good example by the manner in which they embed a spatial correspondence with the real world. The class hierarchies used in many knowledge representation systems constitute a non-spatial analogical representation, with the tree structure of the representation mimicking the hierarchical relation of class inclusion. Analogical representations have long been of interest to the AI community, given their dual abilities to encode information in a perspicuous manner and to facilitate efficient manipulations of that information by exploiting embedded structural properties.

During the past year, we have developed a formal framework for integrating reasoning systems built on analogical and sentential representations; an overview of this work will appear as a paper in the proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (entitled "Reasoning with Analogical Representations"). The framework is comprised of a set of generic

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 – 30 Sep 92

operations on analogical structures and deductive rules for applying those operations. The framework supports both reasoning *about* analogical representations, which amounts to a passive extraction of information from analogical structures for use by a sentential reasoning system, and the more general task of reasoning *with* such representations. The latter casts analogical representations in an active role, having them modified as part of the deductive process. The integration rules were proven sound with respect to an introduced model-theoretic semantics for hybrid systems that combine analogical and sentential representations.

To demonstrate the viability of the formal theory, we implemented a prototype hybrid analogical-sentential reasoner. The implementation was built on top of Mark Stickel's KLAUS automated deduction system, using Myers' technology of universal attachments (described by a forthcoming paper "Hybrid Reasoning using Universal Attachment" to appear in the journal Artificial Intelligence).

Although our analogical-sentential framework was defined independently of any domain, we have explored its application to the problem of reasoning with maps. Our particular focus has been on the type of maps that our mobile robot can generate from perceptual input as it navigates through an office building. Maps built from sensor information generally have gaps corresponding to areas for which perception was unable to determine the relevant physical characteristics, due either to faulty sensors, noise or insufficient perceptual cues. Our hybrid analogical-sentential reasoning framework allows a sentential theory describing properties of the environment to be incorporated into the map-making process. Thus, sentences in a logic can be communicated to the robot as a means of improving upon the information provided by perception alone. This communication provides a means of augmenting sensor-based models of the world with information that is beyond the perceptual capabilities of the robot, leading to more accurate and more complete maps.

Computational methods for reasoning about mental state.

We have continued our work on the development of models of belief, intention, and perception, and logics for reasoning about them. There are three separate research lines: ideal belief systems, representation of intention, and reasoning about causation. Autoepistemic (AE) logic is a formal system characterizing agents that have complete introspective access to their own beliefs. AE logic relies on a fixed point definition that has two significant parts. The first part is a set of assumptions or hypotheses about the contents of the fixed point. The second part is a set of

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 – 30 Sep 92

reflection principles that link sentences with statements about their provability. We have shown, in a paper published in the AAAI conference last July ("Ideal introspective belief") that AE reasoners can be characterized in terms an assumption set of *negative* beliefs about the world (e.g., "I don't believe that I have an older sister"), together with reflection principles relating beliefs to beliefs about beliefs (e.g., "If I believe X, then I believe that I believe X"). We have shown that AE logic is not an ideal logic, in that negative assumptions are too strong for an ideal introspective agent. This theoretical work can help in analyzing metatheoretic systems in logic programming; this further result was presented in an invited paper at the META92 workshop in Uppsala, Sweden ("An autoepistemic analysis of metalevel reasoning in logic programming").

We are also developing a representationalist logic of intention, which we believe is better suited to the properties of intention than the existing normal modal logic of intentions. We are continuing to refine the logic, and, in particular, to introduce principles for intention revision in a paper submitted to the International Joint Conference on AI, "On the Logic of Intention."

We have also continued our work on proof-theoretic techniques for reasoning about mental state, especially on *abduction*. Simply put, abduction is the process of reasoning from some observation to the best explanation for it. Abduction can be used as a reasoning method for many different kinds of problems. Recently we have concentrated on its application to causal and default reasoning, important components of reasoning about mental state. In our previous work, we showed that there are two distinct formalizations for explanatory reasoning. The consistency based approach treats the task as a deductive one, in which the explanation is deduced from a background theory and a minimal set of abnormalities. The abductive method, on the other hand, treats explanations as sentences that, when added to the background theory, derive the observations. We showed that there is a close connection between these two formalizations in the context of simple causal theories: domain theories in which a set of sentences are singled out as the explanatorily relevant causes of observations.

In our current work, we expand the idea of abductive inference in causal theories to include defaults. Our theory is unique in that it integrates a formal notion of causality with nonmonotonic reasoning techniques based on default logic and abduction. The main structure of the theory is a default causal net (DCN) representing the causal connections among propositions in the domain. The causal net provides a framework for the two nonmonotonic reasoning techniques of assuming defaults and generating explanations for observations, allowing them to be combined in a

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 - 30 Sep 92

principled way. This work resulted in a paper to appear in the proceedings of the Third International Conference on the Principles of Knowledge Representation and Reasoning ("A general theory of default reasoning in causal domains").

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 - 30 Sep 92

List of publications, presentations, prizes, and reports

1. **Prizes.**

   In July 1992 we participated in the first AAAI Robotics Exhibition and Competition, which took place in San Jose in conjunction with the AAAI National Conference on AI. We entered Flakey, our mobile robot testbed. Some of the software for this testbed is based on work conducted for this project; specifically, the integration of reaction and planning, and the representation of maps and map-learning. These research areas are mentioned in the summary section on technical progress.

   Flakey performed extremely well in this competition. In the first part, the robots were tested for their ability to react to the environment and stay out of trouble, by not running into moving or stationary obstacles while they roamed around the competition area. Flakey came in second, just one point behind TJ2, the entry from IBM TJ Watson laboratory.

   The second phase of the competition stressed navigation and map-building in real time. The robots were required to explore the arena, locating and recognizing ten poles that were scattered throughout. Having constructed a map, the robots were then asked to visit a sequence of three poles in order.

   Flakey again performed extremely well, coming in second behind CARMEL from the University of Michigan. But there was an important difference between Flakey and all of the other entries. The competition rules allowed each robot team to modify the environment to make the task easier, and every team (except Brown University and us) chose to add visual "bar codes" to the poles to make them easier to locate and identify. We decided to treat the environment as we found it; consequently the exploration task was much harder, and demanded a higher degree of cognitive ability. So, the second-place finish indicates that the perceptual reasoning and map learning techniques that we are developing are a successful means of coping with navigation tasks in a nonengineered environment.

2. **Published and submitted papers.**

   Kurt Konolige and Martha Pollack
   "On the logic of intention."
   Submitted to the International Joint Conference on AI.

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 - 30 Sep 92

Gerd Brewka and Kurt Konolige
"An abductive framework for generalized logic programs."
Submitted to the International Joint Conference on AI.

Kurt Konolige
"What's happening? Elements of commonsense causation."
Proceedings of the International Conference on Cognitive Science, San Sebastian, Spain, May 1991.

Kurt Konolige
"A general theory of default reasoning in causal domains."
Fourth International Workshop on Nonmonotonic Reasoning, Vermont, June 1992.

Kurt Konolige
"Ideal introspective belief."
Proceedings of the National Conference of the American Association for Artificial Intelligence, San Jose, CA, July 1992.

Karen Myers and Kurt Konolige
"Integrating analogical and sentential reasoning for perception."
Proceedings of the AAAI Spring Symposium on Reasoning with Diagrammatic Representations, Stanford, CA, 1992.

Karen Myers and Kurt Konolige
"Semi-autonomous map-making and navigation"
Proceedings of the AAAI Fall Symposium on Applications of AI to Real-world Autonomous Mobile Robots, 1992.

Karen Myers
"Hybrid Reasoning using Universal Attachment"
Artificial Intelligence Journal, to appear

Kurt Konolige
"An autoepistemic analysis of metalevel reasoning in logic programming."
Proceedings of the conference on Metalevel Reasoning and Logic Programming, to be published in Springer Lecture Notes in AI series.

Kurt Konolige and Karen Myers
"Reasoning with analogical representations"
Proceedings of the European Conference on AI Workshop on Theoretical Foundations of Knowledge Representation and Reasoning, Vienna, August 1992

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com   10
Distributed Reasoning and Planning
N00014–89–C–0095
1 Oct 91 – 30 Sep 92

Kurt Konolige
"Using default and causal reasoning in diagnosis"
Proceedings of the European Conference on AI Workshop on Model-based
Reasoning, Vienna, August 1992

Karen Myers and Kurt Konolige
"Reasoning with analogical representations"
Principles of Knowledge Representation and reasoning: Proceedings of the
Third International Conference (KR92), B. Nebel, C. Rich and W. Swartout,
editors, Morgan Kaufmann, 1992., October 1992

Kurt Konolige
"Using default and causal reasoning in diagnosis"
Principles of Knowledge Representation and reasoning: Proceedings of the
Third International Conference (KR92), B. Nebel, C. Rich and W. Swartout,
editors, Morgan Kaufmann, 1992., October 1992

Kurt Konolige, Karen Myers and Alessandro Saffiotti
"CARMEL vs. Flakey: A comparison of two winners"
Submitted to the magazine of the American Association of Artificial Intelli-
gence.

Gerd Brewka, Juergen Dix and Kurt Konolige
"A Tutorial on Nonmonotonic Reasoning"
To be published by in the Center for the Study of Language and Information
Lecture Note Series, Morgan Kaufmann, Los Altos, CA.

3. **Presentations and honors for Kurt Konolige.**

Presentation of papers listed above at various conferences.

Kurt Konolige, Gerd Brewka and Juergen Dix, "A survey of nonmonotonic rea-
soning," invited talk, Nonmonotonic and Inductive Logic Conference, Friedrich-
hoven, Germany, December 1991.

"An autoepistemic analysis of metalevel reasoning in logic programming," in-
vited talk, 3rd International Workshop on Metalevel Reasoning, Uppsala, Swe-
den, June 1992.

(with Gerd Brewka) "Nonmonotonic reasoning," tutorial at the European Con-
ference on AI, Vienna, Austria, August 1992.

Program committee member, ISMIS93, KR92, CADE92, 4th International
Nonmonotonic Workshop, European Workshop on Theoretical Foundations of

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 - 30 Sep 92

Knowledge Representation and Reasoning.

Editorial board, Fundamenta Informaticae, Journal of Applied Non-Classical Logics, International Journal of Applied Intelligence

Reviewer: AIJ, CACM, Journal of Logic Programming

Invited lecturer, Stefan Banach International Mathematical Center, Warsaw, Poland, November 1991.

Co-organizer (with Terry Weymouth), AAAI Robotics Exhibition and Competition, to be held in Washington, DC, July 1993

Invited visiting researcher, NTT Basic Research Laboratories, Tokyo, Japan, September 1992

Co-editor (with Henry Kautz and David Etherington), Special issue of Fundamenta Informaticae on Nonmonotonic Reasoning, To appear in 1993

4. **Presentations and honors for Karen Myers.**

Presentation of papers listed above at various conferences.

"Integrating analogical and sentential reasoning for perception," AAAI Spring Symposium on Reasoning with Diagrammatic Representations, 1992, (Stanford University).

"Attachment methods for integration," AAAI Fall Symposium on Principles of Hybrid Reasoning, 1991 (Asilomar, CA).

Invited panelist for the symposium "Reasoning and visual representations" at the Fourteenth Annual Conference of the Cognitive Science Society, 1992 (Bloomington, Indiana).

"Map-making and Navigation for a Mobile Robot," invited presentation, Center for the Study of Language and Information, Stanford University, May, 1992.

Reviewer for AAAI Magazine, Computational Intelligence

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 - 30 Sep 92

Transitions and DoD interactions

None.

Kurt Konolige
SRI International
(415) 859-2788
konolige@ai.sri.com
Distributed Reasoning and Planning
N00014-89-C-0095
1 Oct 91 – 30 Sep 92

Software and hardware prototypes

Hybrid reasoning: A prototype hybrid reasoning system for combining deductive and analogical reasoning has been constructed. This prototype uses Mark Stickel's KADS theorem prover and Karen Myers' Universal Attachment systems. It is suitable for solving the problems in perceptual reasoning mentioned in the summary of technical progress.

Causal and default reasoning: A prototype abductive reasoning system has been constructed for solving problems in causal and default reasoning. This prototype contains a modification of de Kleer's Assumption-based Truth Maintenance System.

Flakey testbed: We have integrated the work on reactive planning mentioned in the technical summary section into the Flakey mobile robot testbed. The software consists of a fuzzy control subsystem and a simple finite-state automaton task planner. This testbed performed well in the AAAI robotics competition in San Jose (July 1992), capturing 2nd place in both the reactivity and map-making phases of the competition.